

86177-16

Page 53

WE CLAIM:

1. A method of regulating packet flow to a downstream entity capable of forwarding packets to a plurality of intermediate destinations, the method comprising:

10 maintaining a database of queues, each queue in the database being associated with packets intended to be forwarded to a corresponding one of a plurality of final destinations via a corresponding one of the intermediate destinations, each queue in the database being further associated with a state that is either active or inactive;

upon receipt of a message from the downstream entity indicating a reduced ability of a particular one of the intermediate destinations to accept packets intended to be forwarded to a particular one of the final destinations, rendering inactive the state of the queue associated with packets intended to be forwarded to the particular final destination via the particular intermediate destination; and

upon receipt of a message from the downstream entity indicating an increased ability of a particular one of the intermediate destinations to accept packets intended to be forwarded to a particular one of the final destinations, rendering active the state of the queue associated with packets intended to be forwarded to the particular final destination via the particular intermediate destination.

30 2. A method as defined in claim 1, further comprising:
providing storage for the packets associated with
each queue.

3. A method as claimed in claim 2, further comprising:
for each intermediate destination, scheduling
packets for transmission to said intermediate destination
5 from amongst the packets belonging to those queues for
which the state is active and that are associated with
packets intended to be forwarded to any final destination
via said intermediate destination.

10 4. A method as defined in claim 2, further comprising:
determining whether the downstream entity has an
ability to receive at least one packet;
upon determining that the downstream entity has the
ability to receive at least one packet, selecting at
15 least one packet that has been scheduled for transmission
to one of the intermediate destinations and transmitting
at least one of the selected at least one packet to the
downstream entity.

20 5. A method as defined in claim 4, wherein selecting at
least one packet that has been scheduled for transmission
to one of the intermediate destinations includes:
selecting an intermediate destination; and
selecting at least one packet that has been
25 scheduled for transmission to the selected intermediate
destination.

6. A method as defined in claim 1, said method being
implemented in a congestion manager and further
30 comprising:

maintaining information on memory utilization for each of a plurality of flows, each flow being associated with a corresponding one of the final destinations;

5 if memory utilization for a particular one of the flows exceeds a first threshold, generating a message indicative of a reduced ability of the congestion manager to accept packets intended to be forwarded to the final destination associated with the particular flow; and

10 if memory utilization for a particular one of the flows falls below a second threshold, generating a message indicative of an increased ability of the congestion manager to accept packets intended to be forwarded to the final destination associated with the particular flow.

15

7. A method as defined in claim 6, wherein the first and second thresholds are pre-determined.

8. A method as defined in claim 6, further comprising:

20 determining total memory utilization for a plurality of said flows; and

adjusting at least one of the first and second thresholds as a function of the total memory utilization for said plurality of said flows.

25

9. A method as defined in claim 6, further comprising:

receiving the packets from a plurality of upstream entities; and

30 sending the generated messages to the plurality of upstream entities.

PROVISIONAL DRAFT - 6.0

10. A method as defined in claim 9, wherein sending one of the generated messages to the plurality of upstream entities includes broadcasting said one of the generated messages to the plurality of upstream entities.

5

11. A method as defined in claim 9, further comprising: maintaining a second database of those upstream entities that have recently sent packets intended to be forwarded to the final destination associated with the

10 particular flow;

wherein sending the generated messages to the plurality of upstream entities includes multicasting said one of the generated messages to those upstream entities in the second database.

15

12. A method as defined in claim 9, further comprising: maintaining a second database indicative of the number of packets belonging to a given flow that have been received from each upstream entity since 20 transmission to that upstream entity of a message indicative of a reduced ability of the congestion manager to accept packets intended to be forwarded to the final destination associated with the given flow.

25

13. A method as defined in claim 12, further comprising: if the number of packets in the database exceeds a first threshold, for a particular flow and for a particular upstream entity, generating another message indicative of a reduced ability of the congestion manager 30 to accept packets intended to be forwarded to the final destination associated with the particular flow and

PROPRIETARY INFORMATION

sending said other message to the particular upstream entity.

14. A method as defined in claim 13, further comprising:

5 if the number of packets in the database exceeds a second threshold greater than the first threshold, for a particular flow and for a particular upstream entity, generating an alarm message and sending the alarm message to an external entity.

10

15. A method as defined in claim 6, further comprising: maintaining an acknowledgement database, wherein the acknowledgement database includes an entry for each combination of upstream source and final destination and an indication of whether the upstream source in each combination of upstream source and final destination has acknowledged receipt of a message previously sent to the plurality of upstream sources and indicative of an increased ability of the congestion manager to accept 20 packets intended to be forwarded to said final destination.

25

16. A method as defined in claim 15, further comprising: upon receipt of a message from a particular one of the upstream sources acknowledging receipt of a message previously sent to the plurality of upstream sources and indicative of an increased ability of the congestion manager to accept packets intended to be forwarded to a particular final destination, updating the entry in the 30 acknowledgement database corresponding to the combination of particular upstream source and particular final destination.

DRAFT - PROVISIONAL - 08/2001

17. A method as defined in claim 6, further comprising:
upon receipt of a first message from the downstream
entity indicating an increased ability of a particular
5 one of the intermediate destinations to accept packets
intended to be forwarded to the final destination
associated with a particular flow:

generating an acknowledgement message acknowledging
receipt of the first message; and

10 sending the acknowledgement message to the
particular intermediate destination.

18. A method as defined in claim 17, further comprising:
maintaining an acknowledgement database, wherein the
15 acknowledgement database includes an entry for each
combination of upstream source and final destination and
an indication of whether the upstream source in each
combination of upstream source and final destination has
acknowledged receipt of a message previously sent to the
20 plurality of upstream sources and indicative of increased
ability of the congestion manager to accept packets
intended to be forwarded to said final destination.

19. A method as defined in claim 18, further comprising:
25 upon receipt of a message from a particular one of
the upstream sources acknowledging receipt of a message
previously sent to the plurality of upstream sources and
indicative of an increased ability of the congestion
manager to accept packets intended to be forwarded to a
30 particular final destination, updating the entry in the
acknowledgement database corresponding to the combination

PCT/US2001/026660

of particular upstream source and particular final destination.

20. A method as defined in claim 6, further comprising:

5 maintaining an acknowledgement database, wherein the acknowledgement database includes an entry for each combination of upstream source and final destination and an indication of whether the upstream source in each combination of upstream source and final destination has acknowledged receipt of a message previously sent to the plurality of upstream sources and indicative of an increased or reduced ability of the congestion manager to accept packets intended to be forwarded to said final destination.

15

21. A method as defined in claim 20, further comprising:

upon receipt of a message from a particular one of the upstream sources acknowledging receipt of a message previously sent to the plurality of upstream sources and 20 indicative of an increased or reduced ability of the congestion manager to accept packets intended to be forwarded to a particular final destination, updating the entry in the acknowledgement database corresponding to the combination of particular upstream source and 25 particular final destination.

22. A method as defined in claim 6, further comprising:

upon receipt of a first message from the downstream entity indicating an increased or reduced ability of a 30 particular one of the intermediate destinations to accept packets intended to be forwarded to the final destination associated with a particular flow:

DRAFT - DO NOT CITE

generating an acknowledgement message acknowledging receipt of the first message; and

sending the acknowledgement message to the particular intermediate destination.

5

23. A method as defined in claim 21, further comprising:

maintaining an acknowledgement database, wherein the acknowledgement database includes an entry for each combination of upstream source and final destination and an indication of whether the upstream source in each combination of upstream source and final destination has acknowledged receipt of a message previously sent to the plurality of upstream sources and indicative of increased or reduced ability of the congestion manager to accept packets intended to be forwarded to said final destination.

24. A method as defined in claim 23, further comprising:

upon receipt of a message from a particular one of the upstream sources acknowledging receipt of a message previously sent to the plurality of upstream sources and indicative of an increased or reduced ability of the congestion manager to accept packets intended to be forwarded to a particular final destination, updating the entry in the acknowledgement database corresponding to the combination of particular upstream source and particular final destination.

25. A method as defined in claim 6, wherein the congestion manager is implemented in at least one of the intermediate destinations.

86177-16

Page 61

26. A method as defined in claim 1, said method being implemented in at least one of the intermediate destinations.

5 27. A method as claimed in claim 1, wherein the downstream entity is a switch fabric.

10 28. A method as claimed in claim 1, wherein the intermediate destinations include switch fabric nodes of an intermediate switching stage of a multi-stage packet router.

15 29. A method as defined in claim 28, said method being implemented in at least one of the intermediate destinations.

30. A computer-readable storage medium containing program instructions for causing execution in a computing device of a method as defined in claim 1.

20 31. A congestion manager for regulating packet flow to a downstream entity capable of forwarding packets to a plurality of intermediate destinations, comprising:

25 means for maintaining a database of queues, each queue in the database being associated with packets intended to be forwarded to a corresponding one of a plurality of final destinations via a corresponding one of the intermediate destinations, each queue in the database being further associated with a state that is 30 either active or inactive;

means for rendering inactive, upon receipt of a message from the downstream entity indicating a reduced

86177-16

Page 62

ability of a particular one of the intermediate destinations to accept packets intended to be forwarded to a particular one of the final destinations, the state of the queue associated with packets intended to be forwarded to the particular final destination via the particular intermediate destination; and

means for rendering active, upon receipt of a message from the downstream entity indicating an increased ability of a particular one of the intermediate destinations to accept packets intended to be forwarded to a particular one of the final destinations, the state of the queue associated with packets intended to be forwarded to the particular final destination via the particular intermediate destination.

15

32. A computer readable storage medium containing a program element for execution by a computing device to implement a congestion manager for regulating packet flow to a downstream entity capable of forwarding packets to a plurality of intermediate destinations, the program element including:

program code means for maintaining a database of queues, each queue in the database being associated with packets intended to be forwarded to a corresponding one of a plurality of final destinations via a corresponding one of the intermediate destinations, each queue in the database being further associated with a state that is either active or inactive;

program code means for rendering inactive, upon receipt of a message from the downstream entity indicating a reduced ability of a particular one of the intermediate destinations to accept packets intended to

be forwarded to a particular one of the final destinations, the state of the queue associated with packets intended to be forwarded to the particular final destination via the particular intermediate destination;

5 and

program code means for rendering active, upon receipt of a message from the downstream entity indicating an increased ability of a particular one of the intermediate destinations to accept packets intended 10 to be forwarded to a particular one of the final destinations, the state of the queue associated with packets intended to be forwarded to the particular final destination via the particular intermediate destination.

15 33. A congestion manager capable of forwarding packets to a plurality of intermediate destinations, comprising:

a queue processor for maintaining information on a plurality of queues, each queue being associated with packets intended to be forwarded to a corresponding one 20 of a plurality of final destinations via a corresponding one of the intermediate destinations; and

a controller in communication with the queue processor;

said controller being adapted to maintain 25 information on a state of each queue, the state of a queue being either active or inactive;

said controller being further adapted to respond to a message from a particular one of the intermediate destinations indicative of a reduced ability of the 30 particular intermediate destination to accept packets intended to be forwarded to a particular one of the final destinations by rendering inactive the state of the queue

DRAFTED 04-08-2004 - 08:33:00 PM

86177-16

Page 64

associated with packets intended to be forwarded to a particular one of the final destinations via the particular intermediate destination; and

5 said controller being further adapted to respond to a message from a particular one of the intermediate destinations indicative of an increased ability of the particular intermediate destination to accept packets intended to be forwarded to a particular one of the final destinations by rendering active the state of the queue 10 associated with packets intended to be forwarded to a particular one of the final destinations via the particular intermediate destination.

15 34. A congestion manager as defined in claim 33, further comprising:

a memory for providing storage for the packets associated with each queue.

20 35. A congestion manager as defined in claim 34, the queue processor being further adapted to schedule packets for transmission to each particular one of the intermediate destinations from amongst the packets belonging to those queues for which the state is active and that are associated with packets intended to be 25 forwarded to any final destination via the particular intermediate destination.

30 36. A congestion manager as defined in claim 35, the queue processor being further adapted to determine whether the downstream entity has an ability to receive at least one packet and, upon determining that the downstream entity has the ability to receive at least one

86177-16

Page 65

packet, to select at least one packet that has been scheduled for transmission to one of the intermediate destinations and to cause at least one of the selected at least one packet to be extracted from the memory and 5 transmitted to the downstream entity.

37. A congestion manager as defined in claim 33, further comprising:

10 a classifier connected to the queue processor and to the controller, for determining the final destination to which each packet in a received stream of packets is to be forwarded.

15 38. A method of regulating packet flow to a downstream entity capable of forwarding packets to a plurality of intermediate destinations, the method comprising:

20 maintaining a database of queues, each queue in the database being associated with packets of a corresponding one of a plurality of service classes intended to be forwarded to a corresponding one of a plurality of final destinations via a corresponding one of the intermediate destinations, each queue in the database being further associated with a state that is either active or inactive;

25 upon receipt of a message from the downstream entity indicating a reduced ability of a particular one of the intermediate destinations to accept packets of a particular one of the service classes intended to be forwarded to a particular one of the final destinations, 30 rendering inactive the state of the queue associated with packets of the particular service class intended to be

86177-16

Page 66

forwarded to the particular final destination via the particular intermediate destination; and

upon receipt of a message from the downstream entity indicating an increased ability of a particular one of 5 the intermediate destinations to accept packets of a particular one of the service classes intended to be forwarded to a particular one of the final destinations, rendering active the state of the queue associated with packets of the particular service class intended to be 10 forwarded to the particular final destination via the particular intermediate destination.

39. A method as defined in claim 38, further comprising:
providing storage for the packets associated with 15 each queue.

40. A method as claimed in claim 39, further comprising:
for each intermediate destination, scheduling 20 packets for transmission to said intermediate destination from amongst the packets belonging to those queues for which the state is active and that are associated with packets intended to be forwarded to any final destination via said intermediate destination.

25 41. A method as claimed in claim 40, further comprising:
maintaining a queue of active queues for each 30 service class, wherein each queue in the queue of queues for a particular service class has a state that is active; and
for each intermediate destination, scheduling packets for transmission to said intermediate destination

86177-16

Page 67

from amongst the packets in the queues contained in each queue of active queues.

42. A computer-readable storage medium containing
5 program instructions for causing execution in a computing device of a method as defined in claim 38.